Catching Fire via 'Likes': Inferring Topic Preferences of Trump Followers on Twitter

Yu Wang
Political Science
University of Rochester
Rochester, NY 14627
ywang@ur.rochester.edu

Jiebo Luo
Computer Science
University of Rochester
Rochester, NY 14627
jluo@cs.rochester.edu

Richard Niemi
Political Science
University of Rochester
Rochester, NY 14627
niemi@rochester.edu

Yuncheng Li
Computer Science
University of Rochester
Rochester, NY 14627
yli@cs.rochester.edu

Tianran Hu
Computer Science
University of Rochester
Rochester, NY 14627
thu@cs.rochester.edu

arXiv:1603.03099v1 [cs.SI] 9 Mar 2016

Abstract

In this paper, we propose a framework to infer the topic preferences of Donald Trump's followers on Twitter. We first use latent Dirichlet allocation (LDA) to derive the weighted mixture of topics for each Trump tweet. Then we use negative binomial regression to model the "likes," with the weights of each topic serving as explanatory variables. Our study shows that attacking Democrats such as President Obama and former Secretary of State Hillary Clinton earns Trump the most "likes." Our framework of inference is generalizable to the study of other politicians.

Introduction

Republican presidential front-runner Donald Trump has 5.46 million followers on Twitter. Between September 18th and December 27nd, 2015, Trump posted an average of 21.0 tweets per day and each Trump tweet received an average of 3410.7 "likes" with a standard deviation of 2732.7. In this paper, we analyze these variations in "likes."

Figure 1 shows four tweets posted by Donald Trump on December 27, 2015. The first and the fourth of the tweets are about Hillary Clinton, former Secretary of State, who is currently leading in the Democratic presidential race. The second tweet is on Marco Rubio, a presidential candidate from the Republican party.[1] The third tweet is about Donald Trump appearing in the news media. What motivates our study is the large variations in the number of "likes." In this example, they range from 2,200 to 11,000.

We interpret these "likes" as measures of the tweets' attractiveness to Trump's followers: if followers like a Trump tweet, they 'like' it; otherwise, they do not. We further assume that each tweet as a document represents a mixture of topics and that the variations in these "likes" can be partly attributed to these topics, thus revealing the followers' preference. The most preferred topic is expected to earn Trump the most "likes."

[1]For a detailed summary of the many poll results, please see http://elections.huffingtonpost.com/pollster/2016-national-gop-primary.

Figure 1: Examples of Trump's Tweets and the "likes"

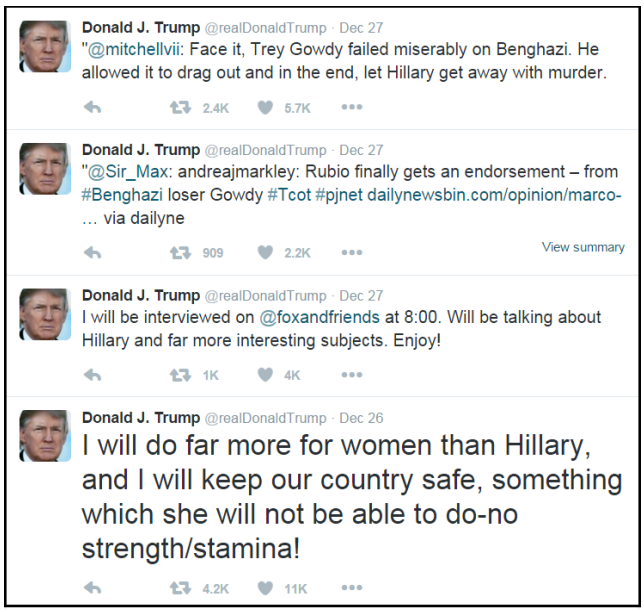
Donald J. Trump @realDonaldTrump · Dec 27
"@mitchellvii: Face it, Trey Gowdy failed miserably on Benghazi. He allowed it to drag out and in the end, let Hillary get away with murder.
2.4K 5.7K
Donald J. Trump @realDonaldTrump · Dec 27
"@Sir_Max: andreajmarkley: Rubio finally gets an endorsement – from #Benghazi loser Gowdy #Tcot #pjnet dailynewsbin.com/opinion/marco-... via dailyne
909 2.2K
View summary
Donald J. Trump @realDonaldTrump · Dec 27
I will be interviewed on @foxandfriends at 8:00. Will be talking about Hillary and far more interesting subjects. Enjoy!
1K 4K
Donald J. Trump @realDonaldTrump · Dec 26
I will do far more for women than Hillary, and I will keep our country safe, something which she will not be able to do-no strength/stamina!
4.2K 11K

To operationalize this idea, we first apply latent Dirichlet allocation (LDA) to extract the topics (Blei, Ng, and Jordan 2003; Chaney and Blei 2012). Then we use negative binomial regression to model the "likes," with the weights of each topic serving as explanatory variables. Throughout, we shall control for other factors that influence the number of "likes": time of the posting, weekend, public debates and the number of Trump followers.

The contributions of our study are threefold. First we contribute to the study of social behavior on Twitter. The existing literature has mostly focused on analyzing the retweeting behavior. Our study complements this literature by exploring the "like" behavior. Second, our study shows that the most favored topic for the Trump followers is attacking the Democrats. This result shall prove valuable for both political scientists and politicians. Third, we propose a framework for inferring followers' topic preference. This framework can be applied to the study of other politicians.

Related Work

Our work builds upon previous research on electoral studies using social media data and on behavioral studies in social media.

There are a large number of studies on using social media data to analyze and forecast election results. (DiGrazia et al. 2013) finds a statistically significant relationship between tweets and electoral outcomes. (MacWilliams 2015) suggests that a candidate's number of "likes" in Facebook can be used for measuring a campaign's success in engaging the public. According to (Williams and Gulati 2008), the number of Facebook fans constitutes an indicator of candidate viability. (Wang, Li, and Luo 2016) uses the profile images on Twitter to study the demographic characteristics of Donald Trump's and Hillary Clinton's followers. Our work uses both the number of Trump followers (as a control variable) and the number of "likes" (as the dependent variable). Our contribution is to infer follower preferences from these "likes."

There are also quite a few studies modeling individual behaviors in social media. (Lee et al. 2015) models the decision to retweet, using Twitter user features such as agreeableness, number of tweets posted, and daily tweeting patterns. (Mahmud, Chen, and Nichols 2013) models individuals' waiting time before replying to a tweet based on their previous replying patterns. Our study models the number of "likes" that a Trump's tweet receives. Our innovation is to use tweet-specific features instead of individual-specific features, as done in the above-cited literature.

Data and Methodology

We use the dataset *US2016*, constructed by us with Twitter data. The dataset contains a tracking record of the number of followers for all the major candidates in the 2016 presidential race, including Donald Trump (Figure 2). The dataset spans the entire period between September 18th, 2015 and December 27th, 2015 and covers three Democratic debates and three Republican debates. In addition, *US2016* also contains all the tweets (2120, in total) that Trump posted during the same period and the number of "likes" that each tweet has received (Figure 3). In Table 1, we report the summary statistics of the main variables.

Figure 2: (Time series) Number of Trump Followers

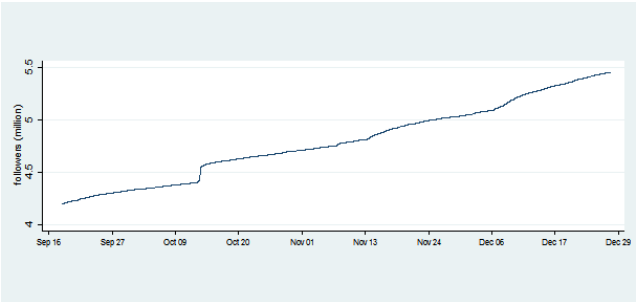

Figure 3: (Time series) 'Likes' per Trump Tweet

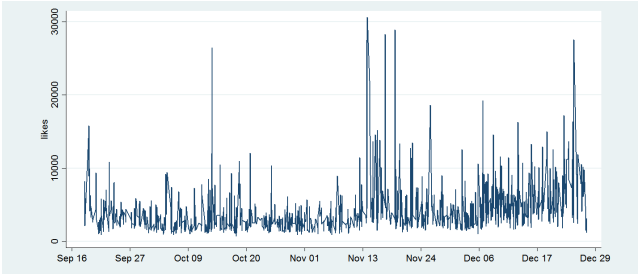

Table 1: Summary statistics

Variable	Min	Max	Mean	S.D.	N
Likes	741	30612	3411	2733	2120
Democratic Debates	0	1	0.088	0.283	2120
Republican Debates	0	1	0.072	0.258	2120
Followers (million)	4.50	5.45	4.788	0.344	2064

LDA and negative binomial regression are the two workhorses of our work. We first use LDA to extract topics from the tweets and then use negative binomial to estimate the coefficients for each topic. In our estimation, we use the following two link functions:

$$\begin{aligned}\mu = exp(&\beta_0 + \beta_1 \text{Weekend} + \beta_2 \text{Democratic Debates} \\ &+ \beta_3 \text{Republican Debates} + \beta_4 \text{Follower Count} \\ &+ \boldsymbol{\gamma} \cdot \textbf{Topic} + \theta \cdot \textbf{Hour Controls}) \\ p = &\, 1/(1+\alpha\mu)\end{aligned}$$

where *Democratic Debates* is binary and takes the value of 1 on the day of a Democratic debate and on the day immediately after the debate, *Republican Debates* is binary and takes the value of 1 on the day of a Republican debate and on the following day. *Weekend* is binary and is 1 if the tweet is posted during the weekend. *Follower Count* is the number of Trump followers when the tweet is posted. **Topic** is a vector and denotes the weights on each topic. **Hour Controls** are a set of dummy variables controlling for the hours when the tweet is posted. α controls the gamma distribution, Gamma($1/\alpha$, α), that generates over-dispersion.

With these two equations, we formulate the likelihood as follows and use it to estimate the coefficients:

$$\frac{\Gamma(1/\alpha + y)}{\Gamma(y+1)\Gamma(1/\alpha)} p^{1/\alpha}(1-p)^y$$

where y denotes the number of "likes."

To select the appropriate number of topics for the LDA procedure, we use as a metric the resulting mean absolute error (MAE) of the negative binomial regression.

In Figure 4, we report the MAE as a function of the number of topics, from 2 to 9. We observe that MAE first decreases when the number of topics increases from 2 to 4 and then remains relatively flat. MAE achieves its minimum

when the number of topics equals 4, so we shall set the number of topics to 4 in the subsequent analysis.[2]

Figure 4: Mean Absolute Error as a Function of Topics

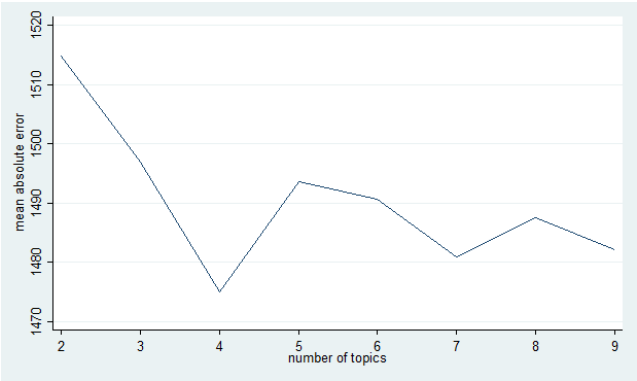

Main Results

In this section, we first present the four topics that we uncover from the LDA procedure. We then report our main estimation results. Lastly, we discuss the estimated differences between these topics.

Topic Classification

In Table 3, we report the four topics that we obtain after applying the LDA procedure to the Trump tweets, together with the top 20 topic words.

We observe these tweets can be classified into four topics: *News Media*, *Republicans*, *Democrats*, and *Trump Campaign*. *News Media* contains such topic words as foxnews, cnn, interview and abc. *Republicans* includes names of Republican candidates such as Jeb Bush, Marco Rubio, Ben Carson and Ted Cruz. *Democrats* includes names from the Democratic party, such as Hillary Clinton and President Obama. *Trump Campaign* relates to Trump's performance in polls and speeches. It includes such words as polls and crowd.

Based on this classification, the first, third, and fourth tweets in Figure 1 shall assign large weights to *Democrats* as they all relate to Hillary Clinton. The second tweet assigns a large weight to the topic *Republicans* as it relates to Marco Rubio, a presidential candidate from the Republican party. And the third tweet assigns a large weight also to the *News Media* topic, as it is about Donald Trump appearing in the news.

Estimation

In this subsection, we use the weights derived from the LDA procedure to estimate the topic preferences of the Trump followers on Twitter. We follow the formulation presented in Section 3 and we report the estimation results in Table 2.

[2]Selecting a much larger number of topics, for example 17, yields an even smaller MAE. But such a larger number of topics will make interpretation very difficult.

Column 1 does not include the extracted topics as explanatory variables and serves as the baseline. It shows that tweets Trump posted during Democratic debates tend to receive more "likes" and tweets posted during Republican debates tend to receive fewer "likes." A similar observation is reported in (Wang, Li, and Luo 2016). The coefficient on *Follower Count* is positive and statistically significant, suggesting that the more followers Trump has the more "likes" a Trump tweet will receive. We also find that tweets posted during the weekend tend to receive fewer "likes."

By contrast, Column 2 includes the topic variables. First, by the Akaike Information Criterion (AIC), we shall select the Column 2 specifications. Second, we exclude the topic *News Media* from the regression to avoid perfect multicollinearity as the four topics sum up to 1. We find that the topics *Democrats* and *Trump Campaign* are statistically more preferred to *News Media*. The difference between *Republicans* and *News Media* is not statistically significant.

Table 2: Negative Binomial Regression

	Baseline	Topics
likes		
Democratic Debates	0.365***	0.286***
	(0.0423)	(0.0412)
Republican Debates	-0.225***	-0.165***
	(0.0479)	(0.0463)
Follower Count	0.691***	0.629***
	(0.0335)	(0.0326)
Weekend	-0.135***	-0.113***
	(0.0269)	(0.0260)
Republicans		0.0368
		(0.0447)
Democrats		0.545***
		(0.0456)
Trump Campaign		0.0840*
		(0.0421)
Constant	4.688***	4.878***
	(0.217)	(0.212)
ln(α)		
Constant	-1.270***	-1.346***
	(0.0298)	(0.0299)
Observations	2063	2063
AIC	36312.5	36148.2

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Using likelihood ratio test on α, we are further able to confirm the existence of over-dispersion, and thus confirm that negative binomial regression is more appropriate than Poisson regression.

The Most Preferred Trump Topic

To visualize the results, we plot the estimated coefficients with 95% confidence intervals in Figure 3. *News Media*, taking the value of 0, serves as the baseline for comparison. The

Table 3: Topic Classification

Topics	Top 20 Topic Words
News Media	trump great america thank again donald make foxnews makeamericagreatagain nice thanks cnn interview president good more tonight job vote need
Republicans	trump cnn poll rubio carson won jeb debate bush jebbush campaign gop marco immigration people foxnews still candidate megynkelly money
Democrats	hillary president why people obama like many going get isis right clinton country want need doing them against foxnews last
Trump Campaign	trump great new poll thank makeamericagreatagain live donald crowd america iowa big night morning book cnn tonight people tomorrow polls

confidence intervals for *Democrats* and *Trump Campaign* are both above 0, suggesting that Trump followers prefer *Democrats* and *Campaign* to *News Media*. By comparison, the confidence interval for *Republicans* is not strictly positive.

More important, we observe that *Democrats* is the most preferred topic among Trump followers, which is consistent with our earlier observation that tweets posted during Democratic debates tend to receive more "likes." Referring back to the quoted tweets in Figure 1, this helps explain why tweets mentioning Hillary Clinton receive more "likes."

Figure 5: Estimated Topic Coefficients (95% C.I.)

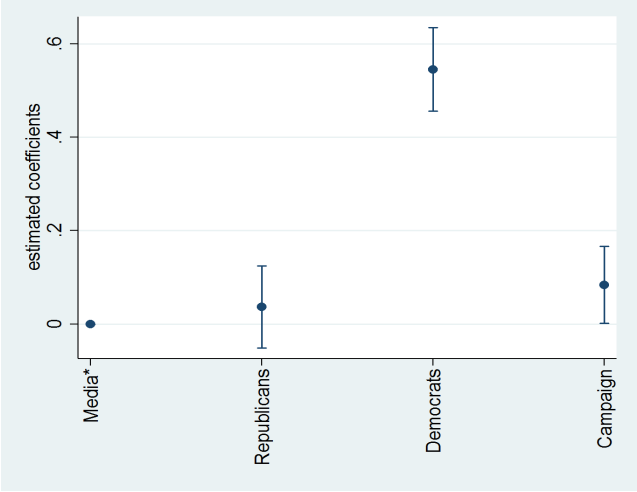

Conclusion and Future Work

We have presented a study of the topic preferences of the presidential candidate Donald Trump's followers on Twitter. We first used latent Dirichlet allocation (LDA) to extract topics from Trump's tweets and then regressed the number of "likes" on the weighted mixture of topics. We used Mean Absolute Error (MAE) from the negative binomial regression as a metric to select the appropriate number of topics. We found that *Democrats* is the most preferred topic among Trump followers. While our work focuses on Donald Trump, the framework of inference that we propose here can be applied to the study of other politicians, such as Hillary Clinton and President Barack Obama.

We believe the rise of Donald Trump is a significant event in American politics. Our immediate next step is to understand the demographics of the Trump followers on Twitter and evaluate their sentiments through, e.g., tweets.

Acknowledgment

We gratefully acknowledge support from the University of Rochester, New York State through the Goergen Institute for Data Science, and our corporate sponsors Xerox and Yahoo.

References

[Blei, Ng, and Jordan 2003] Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research* 3:993–1022.

[Chaney and Blei 2012] Chaney, A. J., and Blei, D. M. 2012. Visualizing topic models. In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media*.

[DiGrazia et al. 2013] DiGrazia, J.; McKelvey, K.; Bollen, J.; and Rojas, F. 2013. More tweets, more votes: Social media as a quantitative indicator of political behavior. *PlOS ONE* 8:1–5.

[Lee et al. 2015] Lee, K.; Mahmud, J.; Chen, J.; and Zhou, M. 2015. Who will retweet this? detecting strangers from twitter to retweet information. *ACM Transactions on Intelligent Systems and Technology* 4:1–25.

[MacWilliams 2015] MacWilliams, M. C. 2015. Forecasting Congressional Elections Using Facebook Data. *PS: Political Science & Politics* 48(04).

[Mahmud, Chen, and Nichols 2013] Mahmud, J.; Chen, J.; and Nichols, J. 2013. When will you answer this? estimating response time in twitter. In *Proceedings of the Seventh International AAAI Conference on Weblogs and Social Media*, 697–700.

[Wang, Li, and Luo 2016] Wang, Y.; Li, Y.; and Luo, J. 2016. Deciphering the 2016 U.S. Presidential Campaign in the Twitter Sphere: A Comparison of the Trumpists and Clintonists. In *Tenth International AAAI Conference on Web and Social Media*.

[Williams and Gulati 2008] Williams, C. B., and Gulati, G. J. 2008. The Political Impact of Facebook: Evidence from the 2006 Midterm Elections and 2008 Nomination Contest. Politics & Technology Review, March, 11-21. *Politics & Technology Review* 11–21.